arXiv:1405.5960v1 [cs.LG] 23 May 2014

LASS: a simple assignment model with Laplacian smoothing

Miguel Á. Carreira-Perpiñán Weiran Wang
Electrical Engineering and Computer Science, University of California, Merced
http://eecs.ucmerced.edu

May 22, 2014

Abstract

We consider the problem of learning soft assignments of N items to K categories given two sources of information: an item-category similarity matrix, which encourages items to be assigned to categories they are similar to (and to not be assigned to categories they are dissimilar to), and an item-item similarity matrix, which encourages similar items to have similar assignments. We propose a simple quadratic programming model that captures this intuition. We give necessary conditions for its solution to be unique, define an out-of-sample mapping, and derive a simple, effective training algorithm based on the alternating direction method of multipliers. The model predicts reasonable assignments from even a few similarity values, and can be seen as a generalization of semisupervised learning. It is particularly useful when items naturally belong to multiple categories, as for example when annotating documents with keywords or pictures with tags, with partially tagged items, or when the categories have complex interrelations (e.g. hierarchical) that are unknown.

1 Introduction

A major success in machine learning in recent years has been the development of semisupervised learning (SSL) (Chapelle et al., 2006), where we are given labels for only a few of the training points. Many SSL approaches rely on a neighborhood graph constructed on the training data (labeled and unlabeled), typically weighted with similarity values. The Laplacian of this graph is used to construct a quadratic nonnegative function that measures the agreement of possible labelings with the graph structure, and minimizing it given the existing labels has the effect of propagating them over the graph. Laplacian-based formulations are conceptually simple, computationally efficient (since the Laplacian is usually sparse), have a solid foundation in graph theory and linear algebra (Chung, 1997; Doyle and Snell, 1984), and most importantly work very well in practice. The graph Laplacian has been widely exploited in machine learning, computer vision and graphics, and other areas: as mentioned, in semisupervised learning, manifold regularization and graph priors (Zhu et al., 2003; Belkin et al., 2006; Zhou et al., 2004) for regression, classification and applications such as supervised image segmentation (Grady, 2006), where one solves a Laplacian-based linear system; in spectral clustering (Shi and Malik, 2000), possibly with constraints (Lu and Carreira-Perpiñán, 2008), and spectral dimensionality reduction (Belkin and Niyogi, 2003) and probabilistic spectral dimensionality reduction (Carreira-Perpiñán and Lu, 2007), where one uses eigenvectors of the Laplacian; in clustering, manifold denoising and surface smoothing (Carreira-Perpiñán, 2006; Wang and Carreira-Perpiñán, 2010; Taubin, 1995), where one iterates products of the data with the Laplacian; etc.

We concern ourselves with assignment problems in a semisupervised learning setting, where we have N items and K categories and we want to find soft assignments of items to categories given some information. This information often takes the form of partial tags or annotations, e.g. for pictures in websites such as Flickr, blog entries, etc. Let us consider a specific example where the items are documents (e.g. papers submitted to this conference) and the categories are keywords. Any given paper will likely be associated to a larger or smaller extent with many keywords, but most authors will tag their papers with only a few of them, usually the most distinctive (although, as we know, there may be other reasons). Thus, few papers will be tagged as "computer science" or "machine learning" because those keywords are perceived as redundant given, say,

"semisupervised learning". However, considered in a larger context (e.g. to include biology papers), such keywords would be valuable. Besides, categories may have various correlations that are unknown to us but that affect the assignments. For example, a hierarchical structure implies that "machine learning" belongs to "computer science" (although it does to "applied maths" to some extent as well). In general, we consider categories as sets having various intersection, inclusion and exclusion relations. Section 6 illustrates this in an example. Finally, it is sometimes practical to tag an item as not associated with a certain category, e.g. "this paper is not about regression" or "this patient does not have fever", particularly if this helps to make it distinctive. In summary, in this type of applications, *it is impractical for an item to be fully labeled over all categories, but it is natural for it to be associated or disassociated with a few categories.* This can be coded with item-category similarity values that are positive or negative, respectively, with the magnitude indicating the degree of association, and zero meaning indifference or ignorance. We call this source of information, which is specific for each item irrespectively of other items, the *wisdom of the expert*.

We also consider another practical source of information. Usually it is easy to construct a similarity of a given item to other items, at least its nearest neighbors. For example, with documents or images, this could be based on a bag-of-words representation. We would expect similar items to have similar assignment vectors, and this can be captured with an item-item similarity matrix and its graph Laplacian. We call this source of information, which is about an item in the context of other items, the *wisdom of the crowd*.

As an example of the interaction of these two sources of information, imagine the following example where the items are conference papers and the categories are authors. We know that author A1 writes papers about regression (usually with author A2) or bioinformatics (usually with author A3). We are sent for review a paper that is about regression (it contains many words about regression) and, we are tipped, by A1. Can we guess its coauthors, i.e., predict its assignments to all existing authors? Based on the crowd wisdom alone, many authors could have written the paper (those who write regression papers). Based on the expert wisdom alone, A2 or A3 may have cowritten the paper (since each of them has coathored other papers with A1). Given both wisdoms, we might expect a high assignment for A2 (and A1) and low for everybody else.

In this paper, we propose a simple model that captures this intuition as a quadratic program. We give some properties of the solution, define an out-of-sample mapping, derive a training algorithm, and illustrate the model with document and image datasets.

A shorter version of this work appears in a conference paper (Carreira-Perpiñán and Wang, 2014).

2 The Laplacian assignment (LASS) model

We consider the following assignment problem. We have N items and K categories, and we want to determine soft assignments z_{nk} of each item n to each category k, where $z_{nk} \in [0,1]$ and $\sum_{k=1}^{K} z_{nk} = 1$ for each $n = 1, \dots, N$. We are given two similarity matrices, suitably defined, and typically sparse: an item-item similarity matrix $\mathbf{W}$, which is an $N \times N$ matrix of affinities $w_{nm} \geq 0$ between each pair of items n and m; and an item-category similarity matrix $\mathbf{G}$, which is an $N \times K$ matrix of affinities $g_{nk} \in \mathbb{R}$ between each pair of item n and category k (negative affinities, i.e., dissimilarities, are allowed in $\mathbf{G}$).

We want to assign items to categories optimally as follows:

$$\min_{\mathbf{Z}} \quad \lambda \operatorname{tr}\left(\mathbf{Z}^T \mathbf{L} \mathbf{Z}\right) - \operatorname{tr}\left(\mathbf{G}^T \mathbf{Z}\right) \tag{1a}$$

$$\text{s.t.} \quad \mathbf{Z}\mathbf{1}_K = \mathbf{1}_N \tag{1b}$$

$$\mathbf{Z} \geq \mathbf{0} \tag{1c}$$

where $\lambda > 0$, $\mathbf{1}_K$ is a vector of K ones, and $\mathbf{L}$ is the $N \times N$ graph Laplacian matrix, obtained as $\mathbf{L} = \mathbf{D} - \mathbf{W}$, where $\mathbf{D} = \operatorname{diag}\left(\sum_{n=1}^{N} w_{nm}\right)$ is the degree matrix of the weighted graph defined by $\mathbf{W}$. The problem is a quadratic program (QP) over an $N \times K$ matrix $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_K)$, i.e., NK variables[1] altogether, where $\mathbf{z}_k$, the kth column of $\mathbf{Z}$, contains the assignments of each item to category k. We will call problems of the type (1) *Laplacian assignment problems (LASS)*. Minimizing objective (1a) encourages items to be assigned

[1] We will use a boldface vector $\mathbf{z}_n$ or to mean the nth row of $\mathbf{Z}$ (as a column vector), and a boldface vector $\mathbf{z}_k$ to mean the kth column of $\mathbf{Z}$ (likewise for $\mathbf{g}_n$ or $\mathbf{g}_k$). The context and the index (n or k) will determine which is the case. This mild notation abuse will simplify the explanations.

to categories with which they have high similarity (the linear term in $\mathbf{G}$), while encouraging similar items to have similar assignments (the Laplacian term in $\mathbf{L}$), since

$$\operatorname{tr}\left(\mathbf{Z}^T\mathbf{L}\mathbf{Z}\right) = \frac{1}{2}\sum_{n,m=1}^{N} w_{nm}\left\|\mathbf{z}_n - \mathbf{z}_m\right\|^2$$

where $\mathbf{z}_n$ is the nth row of $\mathbf{Z}$, i.e., the assignments for item n. Although we could absorb λ inside $\mathbf{G}$, we will find it more convenient to fix the scale of each similarity g_{nk} to, say, the interval $[-1, 1]$ (where ± 1 mean maximum (dis)similarity and 0 ignorance), and then let λ control the strength of the Laplacian term.

The objective function (1a) is separable over categories as

$$\sum_{k=1}^{K} E(\mathbf{z}_k; \mathbf{L}, \mathbf{g}_k) \qquad \text{where} \qquad E(\mathbf{z}; \mathbf{L}, \mathbf{g}) = \lambda \mathbf{z}^T\mathbf{L}\mathbf{z} - \mathbf{g}^T\mathbf{z}$$

and the constraints (1b)–(1c) are separable over items as $\mathbf{z}_n^T\mathbf{1}_K = \mathbf{1}_N$, $\mathbf{z}_n \geq \mathbf{0}$, for $n = 1, \ldots, N$. Thus, the problem (1) is not separable, since all the assignments are coupled with a certain structure.

2.1 Extreme values of λ

We can determine the solution(s) for the following extreme values of λ:

- If $\lambda = 0$, then the LASS problem is a linear program (LP) and separates over each item $n = 1, \ldots, N$. The solution is $z_{nk} = \delta(k - k_{\max}(n))$ where $k_{\max}(n) = \arg\max\{g_{nk},\ k = 1, \ldots, K\}$, i.e., each item is assigned to its most similar category. This tells us what the linear term can do by itself. (If the maximum is achieved for more than one category for a given point, then there is an infinite number of "mixed" solutions that correspond to giving any assignment value to those categories, and zero for the rest.)

- If $\lambda = \infty$ or equivalently $\mathbf{G} = \mathbf{0}$, then the LASS problem is a quadratic program with an infinite number of solutions of the form $\mathbf{z}_n = \mathbf{z}_m$ for each $n, m = 1, \ldots, N$, i.e., all items have the same assignments. This tells us what the Laplacian term can do by itself.

- If $\lambda \to \infty$, i.e., for very large λ but still having the linear term, the behavior actually differs from that of $\mathbf{G} = \mathbf{0}$. In the generic case, we expect a unique solution close to $\mathbf{Z} = \mathbf{1}_N\mathbf{z}^T$ where $\mathbf{z} \in \mathbb{R}^K$ and $z_k = \delta(k - k_{\max})$ where $k_{\max} = \arg\max_k(\mathbf{G}^T\mathbf{1}_N)$, $k = 1, \ldots, K$, i.e., all items are assigned to the same category, the one having maximum total similarity $\mathbf{g}_k^T\mathbf{1}_N$ over all items. (Again, if $\mathbf{G}^T\mathbf{1}_N$ has more than one maximum values, there is an infinite number of solutions corresponding to mixed assignments.) Indeed, if λ is very large, the Laplacian term dominates and we have that $\mathbf{z}_n = \mathbf{z}_m = \mathbf{z}$ for every pair of items, approximately. Then the LASS problem becomes the following LP

$$\begin{aligned} \max_{\mathbf{z}} \quad & \mathbf{z}^T(\mathbf{G}^T\mathbf{1}_N) \\ \text{s.t.} \quad & \mathbf{z}^T\mathbf{1}_K = 1 \\ & \mathbf{z} \geq \mathbf{0} \end{aligned}$$

  whose solution allocates all the assignment mass to the category with largest total similarity $\mathbf{g}_k^T\mathbf{1}_N$.

With intermediate $\lambda > 0$, more interesting solutions appear (particularly when the similarity matrices are sparse), where the item-category similarities are propagated to all points through the item-item similarities.

2.2 Existence and unicity of the solution

The LASS problem is a convex QP, so general results of convex optimization tell us that all minima are global minima. However, since the Hessian of the objective function is positive semidefinite, there can be multiple minima. The following theorem characterizes the solutions, and its corollary gives a sufficient condition for the minimum to be unique.

Theorem 2.1. *Assume the graph Laplacian* $\mathbf{L}$ *corresponds to a connected graph and let* $\mathbf{Z}^* \in \mathbb{R}^{N\times K}$ *be a solution (minimizer) of the LASS problem* (1)*. Then, any other solution has the form* $\mathbf{Z}^* + \mathbf{1}_N\mathbf{p}^T$ *where* $\mathbf{p} \in \mathbb{R}^K$ *satisfies the conditions:*

$$\mathbf{p}^T\mathbf{1}_K = 0, \qquad \mathbf{p}^T(\mathbf{G}^T\mathbf{1}_N) = 0, \qquad \mathbf{Z}^* + \mathbf{1}_N\mathbf{p}^T \geq \mathbf{0}. \tag{2}$$

In particular, for each $k = 1, \ldots, K$ *for which* $\exists n \in \{1, \ldots, N\}$: $z^*_{nk} = 0$*, then* $p_k \geq 0$*.*

Proof. Call $q(\mathbf{Z}) = \lambda \operatorname{tr}\left(\mathbf{Z}^T\mathbf{L}\mathbf{Z}\right) - \operatorname{tr}\left(\mathbf{G}^T\mathbf{Z}\right)$ the objective function of the LASS problem. Since q is continuous and the feasible set of the problem is bounded and closed in $\mathbb{R}^{N\times K}$, q achieves a minimum value in the feasible set, hence at least one solution exists, which makes the theorem statement well defined. We call this solution $\mathbf{Z}^*$. Now let us show that $q(\mathbf{Z}^* + \mathbf{P}) \geq q(\mathbf{Z}^*)$ for any other feasible point $\mathbf{Z}^* + \mathbf{P}$, with $\mathbf{P} \in \mathbb{R}^{N\times K}$. Simple algebra shows that

$$q(\mathbf{Z}^* + \mathbf{P}) = q(\mathbf{Z}^*) + \operatorname{tr}\left(\mathbf{P}^T(2\lambda\mathbf{L}\mathbf{Z}^* - \mathbf{G})\right) + \lambda \operatorname{tr}\left(\mathbf{P}^T\mathbf{L}\mathbf{P}\right). \tag{3}$$

The last term is nonnegative because $\mathbf{L}$ is positive semidefinite. The penultimate term is also nonnegative. To see this, write the KKT conditions for the LASS problem with Lagrange multipliers $\boldsymbol{\pi} \in \mathbb{R}^N$ and $\mathbf{M} \in \mathbb{R}^{N\times K} = (\mu_{nk})$ for the equality and inequality constraints, respectively:

$$\begin{aligned} 2\lambda\mathbf{L}\mathbf{Z}^* - \mathbf{G} - \boldsymbol{\pi}\mathbf{1}_K^T - \mathbf{M} &= \mathbf{0} \\ \mathbf{Z}^*\mathbf{1}_K &= \mathbf{1}_N \\ \mathbf{Z}^* &\geq \mathbf{0} \\ \mathbf{M} &\geq \mathbf{0} \\ \mu_{nk}z^*_{nk} = 0,\ n = 1, \ldots, N,\ k = 1, \ldots, K. & \end{aligned}$$

Thus, since $\mathbf{Z}^* + \mathbf{P}$ is feasible, $(\mathbf{Z}^* + \mathbf{P})\mathbf{1}_K = \mathbf{1}_N \Rightarrow \mathbf{P}\mathbf{1}_K = \mathbf{0}$, and $\mathbf{Z}^* + \mathbf{P} \geq \mathbf{0} \Rightarrow p_{nk} \geq 0$ where $z^*_{nk} = 0$ (i.e., the active inequalities). Then, from the first KKT equation, we have for the penultimate term:

$$\operatorname{tr}\left(\mathbf{P}^T(2\lambda\mathbf{L}\mathbf{Z}^* - \mathbf{G})\right) = \operatorname{tr}\left(\mathbf{P}^T\boldsymbol{\pi}\mathbf{1}_K^T + \mathbf{P}^T\mathbf{M}\right) = \operatorname{tr}\left(\mathbf{1}_K^T\mathbf{P}^T\boldsymbol{\pi} + \mathbf{P}^T\mathbf{M}\right) = \sum_{n,k=1}^{N,K} p_{nk}\mu_{nk} = \sum_{\text{active } n,k} p_{nk}\mu_{nk} \geq 0$$

because $\mu_{nk} \geq 0$ for all n, k, $\mu_{nk} = 0$ for the inactive inequalities, and $p_{nk} \geq 0$ for the active inequalities. Hence, the last two terms in (3) are nonnegative, so $q(\mathbf{Z}^* + \mathbf{P}) \geq q(\mathbf{Z}^*)$ and $\mathbf{Z}^*$ is a global minimizer.

Now assume that $q(\mathbf{Z}^* + \mathbf{P}) = q(\mathbf{Z}^*)$. Then the last two terms in (3) must both be zero. Recall that if the graph Laplacian $\mathbf{L}$ corresponds to a connected graph, it has a single null eigenvalue with an eigenvector consisting of all ones. From $\operatorname{tr}\left(\mathbf{P}^T\mathbf{L}\mathbf{P}\right) = 0$ it follows that $\mathbf{P} = \mathbf{1}_N\mathbf{p}^T$ for some $\mathbf{p} \in \mathbb{R}^K$. Since $\mathbf{Z}^* + \mathbf{P}$ is feasible, $\mathbf{P}\mathbf{1}_K = \mathbf{0} \Rightarrow \mathbf{p}^T\mathbf{1}_K = 0$, and $\mathbf{Z}^* + \mathbf{1}_N\mathbf{p}^T \geq \mathbf{0}$. From the penultimate term, $0 = \operatorname{tr}\left(\mathbf{P}^T(2\lambda\mathbf{L}\mathbf{Z}^* - \mathbf{G})\right) = \operatorname{tr}\left(2\lambda\mathbf{p}\mathbf{1}_N^T\mathbf{L}\mathbf{Z}^* - \mathbf{p}\mathbf{1}_N^T\mathbf{G}\right) \Rightarrow \mathbf{p}^T(\mathbf{G}^T\mathbf{1}_N) = 0$.

Finally, from $\mathbf{Z}^* + \mathbf{1}_N\mathbf{p}^T \geq \mathbf{0}$ it follows that $z^*_{nk} + p_{nk} \geq 0$ for each $n = 1, \ldots, N$ and $k = 1, \ldots, K$, so if $\exists n, k$: $z^*_{nk} = 0$ (i.e., if any of the inequalities involving k are active), then $p_k \geq 0$. □

Corollary 2.2. *Assume the graph Laplacian* $\mathbf{L}$ *corresponds to a connected graph and let* $\mathbf{Z}^* \in \mathbb{R}^{N\times K}$ *be a solution of the LASS problem* (1)*. If* $\max_k\left(\min_n\left(z^*_{nk}\right)\right) = 0$ *then the solution* $\mathbf{Z}^*$ *is unique.*

Proof. We have $\max_k\left(\min_n\left(z^*_{nk}\right)\right) = 0$, so for each k, $\exists n$: $z^*_{nk} = 0$. From theorem 2.1, any other solution $\mathbf{Z}^* + \mathbf{1}_N\mathbf{p}^T$ must have $p_k \geq 0$ for $k = 1, \ldots, K$, and $\mathbf{p}^T\mathbf{1}_K = 0$. Hence $\mathbf{p} = \mathbf{0}$. □

Remark 2.3. If the graph Laplacian $\mathbf{L}$ corresponds to a graph with multiple connected components, then the LASS problem separates into a problem for each component, and the previous theorem holds in each component. Computationally, it is also more efficient to solve each problem separately.

Remark 2.4. The set (2) of solutions to a LASS problem is a convex polytope.

Remark 2.5. The condition of corollary 2.2 means that each category has at least one item with a zero assignment to it. In practice, we can expect this condition to hold, and therefore the solution to be unique, if the categories are sufficiently distinctive and λ is small enough. Equivalently, nonunique solutions arise if some categories are a catch-all for the entire dataset. Theoretically, this should always be possible if λ is large

enough, particularly if there are many categories. However, in practice we have never observed nonunique solutions, because for large λ the algorithm is attracted towards a solution where one category dominates, so the vector $\mathbf{p}$ can only have one negative component, which makes (2) impossible unless $\mathbf{G}$ takes a special value, such as $\mathbf{G} = \mathbf{0}$. Thus, the symmetric situation where all assignments are possible for $\lambda \to \infty$ does not seem to occur in practice.

Remark 2.6. Practically, one can always make the solution unique by replacing $\mathbf{L}$ with $\mathbf{L} + \epsilon \mathbf{I}_N$ where $\epsilon > 0$ is a small value, since this makes the objective strongly convex. (This has the equivalent meaning of adding a penalty $\epsilon \|\mathbf{Z}\|_F^2$ to it, which has the effect of biasing the assignment vector $\mathbf{z}_n$ of each item towards the simplex barycenter, i.e., uniform assignments.) However, as noted before, nonunique solutions appear to be rare with practical data, so this does not seem necessary.

2.3 Particular cases

Theorem 2.7. *Assume the graph Laplacian* $\mathbf{L}$ *corresponds to a connected graph and let* $\mathbf{Z} \in \mathbb{R}^{N \times K}$ *be a solution of the LASS problem* (1)*. Then:*

1. *If* $\mathbf{g}_n = \mathbf{0}$ *then* $\mathbf{z}_n = \frac{\mathbf{Z}^T \mathbf{w}_n}{\mathbf{1}_N^T \mathbf{w}_n} = \sum_{m=1}^N \frac{w_{nm}}{\sum_{m'=1}^N w_{nm'}} \mathbf{z}_m$, $\boldsymbol{\mu}_n = \mathbf{0}$ *and* $\pi_n = 0$.
2. *If* $\mathbf{g}_k \leq \mathbf{0}$ *then* $\mathbf{z}_k = \mathbf{0}$, $\boldsymbol{\mu}_n = -\mathbf{g}_k$ *and* $\pi_n = 0$.

Proof. Both statements follow from substituting the values in the KKT conditions (4). Conditions (4b)–(4e) are trivially satisfied, so we prove condition (4a) only. For statement 1, we can write row n of $\mathbf{L}$ as $\mathbf{l}_n = -\mathbf{w}_n + (\mathbf{1}_N^T \mathbf{w}_n)\mathbf{e}_n$, where $\mathbf{w}_n$ is row n of $\mathbf{W}$ and $\mathbf{e}_n$ is a vector with entries $e_{nm} = \delta(n - m)$, and we write all row vectors as column vectors. Then we can write row n (in column form) of condition (4a) as:

$$\mathbf{0} = 2\lambda \mathbf{Z}^T \mathbf{l}_n - \mathbf{g}_n - \pi_n \mathbf{1}_K - \boldsymbol{\mu}_n = 2\lambda(-\mathbf{Z}^T \mathbf{w}_n + (\mathbf{1}_N^T \mathbf{w}_n)\mathbf{z}_n) \Rightarrow \mathbf{z}_n = \frac{\mathbf{Z}^T \mathbf{w}_n}{\mathbf{1}_N^T \mathbf{w}_n}.$$

For statement 2, we can write column k of condition (4a) as $\mathbf{0} = 2\lambda \mathbf{L}\mathbf{z}_k - \mathbf{g}_k - \boldsymbol{\pi} - \boldsymbol{\mu}_k \Rightarrow \boldsymbol{\mu} = -\mathbf{g}_k$. □

Remark 2.8. The meaning of th. 2.7 is as follows. (1) An item for which no similarity to any category is given (i.e., no expert information) receives as assignment the average of its neighbors. This corresponds to the SSL prediction. (2) A category for which no item has a positive similarity receives no assignments.

2.4 Lagrange multipliers for a solution

Given a feasible point $\mathbf{Z}_{N \times K}$ in parameter space, we may want to test whether it is a solution of the LASS problem. For a QP, the KKT conditions are necessary and sufficient for a solution (Nocedal and Wright, 2006). For our problem, and written in matrix form, the KKT conditions are:

$$2\lambda \mathbf{L}\mathbf{Z} - \mathbf{G} - \boldsymbol{\pi}\mathbf{1}_K^T - \mathbf{M} = \mathbf{0} \tag{4a}$$
$$\mathbf{Z}\mathbf{1}_K = \mathbf{1}_N \tag{4b}$$
$$\mathbf{Z} \geq \mathbf{0} \tag{4c}$$
$$\mathbf{M} \geq \mathbf{0} \tag{4d}$$
$$\mathbf{M} \circ \mathbf{Z} = \mathbf{0} \tag{4e}$$

where $\circ$ means elementwise product, and $\boldsymbol{\pi}$ and $\mathbf{M}$ are the Lagrange multipliers associated with the point $\mathbf{Z}$ for the equality and inequality constraints, respectively. We need to compute $\boldsymbol{\pi}$ and $\mathbf{M}$. Given $\mathbf{Z}$, the KKT system (4) has $2NK$ linear equations for $N + NK$ unknowns, and its solution is unique if $\mathbf{Z}$ is feasible, as we will see. To obtain it, we multiply (4a) times $\mathbf{1}_K$ on the right and obtain $\boldsymbol{\pi}$ as a function of $\mathbf{M}$:

$$\boldsymbol{\pi} = \frac{1}{K}(2\lambda \mathbf{L}\mathbf{Z} - \mathbf{G} - \mathbf{M})\mathbf{1}_K = -\frac{1}{K}(\mathbf{G} + \mathbf{M})\mathbf{1}_K. \tag{5}$$

Substituting it in (4a) gives, together with (4e):

$$\mathbf{M} = \mathbf{Q} + \frac{1}{K}\mathbf{M}\mathbf{1}_K\mathbf{1}_K^T$$
$$\mathbf{M} \circ \mathbf{Z} = \mathbf{0}$$

where $\mathbf{Q} = 2\lambda\mathbf{L}\mathbf{Z} - \mathbf{G} + \frac{1}{K}\mathbf{G}\mathbf{1}_K\mathbf{1}_K^T$. This is a linear system of $2NK$ equations for the NK unknowns in $\mathbf{M}$. It separates over each row of $\mathbf{M}$, $n = 1, \dots, N$, in a system of the form

$$\begin{array}{l}\boldsymbol{\mu} - \frac{1}{K}(\mathbf{1}_K^T\boldsymbol{\mu})\mathbf{1}_K = \mathbf{q} \\ \operatorname{diag}(\boldsymbol{\mu})\operatorname{diag}(\mathbf{z}) = \mathbf{0}\end{array} \quad \Longleftrightarrow \quad \begin{pmatrix}\mathbf{I}_K - \frac{1}{K}\mathbf{1}_K^T\mathbf{1}_K \\ \operatorname{diag}(\mathbf{z})\end{pmatrix}\boldsymbol{\mu} = \begin{pmatrix}\mathbf{q} \\ \mathbf{0}\end{pmatrix}$$

where $\boldsymbol{\mu}$, $\mathbf{z}$ and $\mathbf{q}$ correspond to the nth row of $\mathbf{M}$, $\mathbf{Z}$ and $\mathbf{Q}$, respectively, written as K-dimensional column vectors (we omit the index n). This is a linear system of $2K$ equations for K unknowns, which we solve by multiplying on the left times the transpose of the coefficient matrix:

$$\left(\boldsymbol{\Psi} - \frac{1}{K}\mathbf{1}_K\mathbf{1}_K^T\right)\boldsymbol{\mu} = \mathbf{q} - \frac{1}{K}(\mathbf{1}_K^T\mathbf{q})\mathbf{1}_K \qquad \boldsymbol{\Psi} = \operatorname{diag}(\mathbf{z})^2 + \mathbf{I}_K.$$

Finally, using the matrix inversion lemma

$$(\mathbf{A} + \alpha\mathbf{u}\mathbf{u}^T)^{-1} = \mathbf{A}^{-1} - \frac{\alpha\mathbf{A}^{-1}\mathbf{u}\mathbf{u}^T\mathbf{A}^{-1}}{1 + \alpha\mathbf{u}^T\mathbf{A}^{-1}\mathbf{u}}$$

we obtain the solution for $\boldsymbol{\mu}$:

$$\boldsymbol{\mu} = \boldsymbol{\Psi}^{-1}\mathbf{q} - \left(\frac{\mathbf{1}_K^T\mathbf{q} - \mathbf{1}_K^T\boldsymbol{\Psi}^{-1}\mathbf{q}}{K - \mathbf{1}_K^T\boldsymbol{\Psi}^{-1}\mathbf{1}_K}\right)\boldsymbol{\Psi}^{-1}\mathbf{1}_K \tag{6}$$

whose transpose gives row n of $\mathbf{M}$. Note the formula is well defined because $\psi_{kk} = z_k^2 + 1 \geq 0$ and $\mathbf{1}_K^T\boldsymbol{\Psi}^{-1}\mathbf{1}_K = \sum_{k=1}^K \frac{1}{1+z_k^2} < K$ since $z_k \geq 0$ for $k = 1, \dots, K$ and $z_k > 0$ for at least one $k \in \{1, \dots, K\}$, since $\mathbf{Z}$ is feasible.

For the case $\lambda = 0$, one can verify that the above formulas simplify as follows (again, we give $\mathbf{z}$, $\boldsymbol{\mu}$ and π for item n but omitting the index n):

$$\begin{aligned}&\mathbf{z} = (z_k),\ z_k = \delta(k - k_{\max}) \\ &\boldsymbol{\mu} = g_{\max}\mathbf{1}_K - \mathbf{g} \\ &\pi = -g_{\max}\end{aligned}$$

where $\mathbf{g}$ represents the nth row of $\mathbf{G}$, $k_{\max} = \arg\max\{g_k,\ k = 1, \dots, K\}$ and $g_{\max} = g_{k_{\max}}$.

3 A simple, efficient algorithm to solve the QP

It is possible to solve problem (1) in different ways, but one must be careful in designing an effective algorithm because the number of variables and the number of constraints grows proportionally to the number of data points, and can then be very large. We describe here one algorithm that is very simple, has guaranteed convergence without line searches, and takes advantage of the structure of the problem and the sparsity of $\mathbf{L}$. It is based on the alternating direction method of multipliers (ADMM), combined with a direct linear solver using the Schur's complement and caching the Cholesky factorization of $\mathbf{L}$.

3.1 QP solution using ADMM

We briefly review how to solve a QP using the alternating direction method of multipliers (ADMM), following (Boyd et al., 2011). Consider the QP

$$\min_{\mathbf{x}} \quad \frac{1}{2}\mathbf{x}^T\mathbf{P}\mathbf{x} + \mathbf{q}^T\mathbf{x} \tag{7}$$
$$\text{s.t.} \quad \mathbf{A}\mathbf{x} = \mathbf{b},\ \mathbf{x} \geq \mathbf{0} \tag{8}$$

over $\mathbf{x} \in \mathbb{R}^D$, where $\mathbf{P}$ is positive (semi)definite. In ADMM, we introduce new variables $\mathbf{z} \in \mathbb{R}^D$ so that we replace the inequalities with an indicator function $g(\mathbf{z})$, which is zero in the nonnegative orthant $\mathbf{z} \geq \mathbf{0}$ and ∞ otherwise. Then we write the problem as

$$\min_{\mathbf{x}} \quad f(\mathbf{x}) + g(\mathbf{z}) \tag{9}$$

$$\text{s.t.} \quad \mathbf{x} = \mathbf{z} \tag{10}$$

where

$$f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T\mathbf{P}\mathbf{x} + \mathbf{q}^T\mathbf{x}, \qquad \text{dom}\,(f) = \{\mathbf{x} \in \mathbb{R}^D \colon \mathbf{A}\mathbf{x} = \mathbf{b}\}$$

is the original objective with its domain restricted to the equality constraint. The augmented Lagrangian is

$$\mathcal{L}(\mathbf{x}, \mathbf{z}, \mathbf{y}; \rho) = f(\mathbf{x}) + g(\mathbf{z}) + \mathbf{y}^T(\mathbf{x} - \mathbf{z}) + \frac{\rho}{2}\|\mathbf{x} - \mathbf{z}\|^2 \tag{11}$$

and the ADMM iteration has the form:

$$\begin{aligned}
\mathbf{x} &\leftarrow \arg\min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \mathbf{z}, \mathbf{y}; \rho) \\
\mathbf{z} &\leftarrow \arg\min_{\mathbf{z}} \mathcal{L}(\mathbf{x}, \mathbf{z}, \mathbf{y}; \rho) \\
\mathbf{y} &\leftarrow \mathbf{y} + \rho(\mathbf{x} - \mathbf{z})
\end{aligned}$$

where $\mathbf{y}$ is the dual variable (the Lagrange multiplier estimates for the constraint $\mathbf{x} = \mathbf{z}$), and the updates are applied in order and modify the variables immediately. Here, we use the scaled form of the ADMM iteration, which is simpler. It is obtained by combining the linear and quadratic terms in $\mathbf{x} - \mathbf{z}$ and using a scaled dual variable $\mathbf{u} = \mathbf{y}/\rho$:

$$\begin{aligned}
\mathbf{x} &\leftarrow \arg\min_{\mathbf{x}} \left(f(\mathbf{x}) + \frac{\rho}{2}\|\mathbf{x} - \mathbf{z} + \mathbf{u}\|^2\right) \\
\mathbf{z} &\leftarrow \arg\min_{\mathbf{z}} \left(g(\mathbf{z}) + \frac{\rho}{2}\|\mathbf{x} - \mathbf{z} + \mathbf{u}\|^2\right) \\
\mathbf{u} &\leftarrow \mathbf{u} + \mathbf{x} - \mathbf{z}.
\end{aligned}$$

Since $g(\mathbf{x})$ is the indicator function for the nonnegative orthant, the solution of the $\mathbf{z}$-update is simply to threshold each entry in $\mathbf{x} + \mathbf{u}$ by taking is nonnegative part. Finally, the ADMM iteration is:

$$\mathbf{x} \leftarrow \arg\min_{\mathbf{x}} \left(f(\mathbf{x}) + \frac{\rho}{2}\|\mathbf{x} - \mathbf{z} + \mathbf{u}\|^2\right) \tag{12}$$

$$\mathbf{z} \leftarrow (\mathbf{x} + \mathbf{u})_+ \tag{13}$$

$$\mathbf{u} \leftarrow \mathbf{u} + \mathbf{x} - \mathbf{z} \tag{14}$$

where the updates are applied in order and modify the variables immediately, and $(t)_+ = \max(t, 0)$ applies elementwise, and $\|\cdot\|$ is the Euclidean norm. The penalty parameter $\rho > 0$ is set by the user, and $\mathbf{z} = (z_1, \ldots, z_D)^T$ are the Lagrange multiplier estimates for the inequalities. The $\mathbf{x}$-update is an equality-constrained QP with KKT conditions

$$\begin{pmatrix} \mathbf{P} + \rho\mathbf{I} & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \boldsymbol{\nu} \end{pmatrix} = \begin{pmatrix} -\mathbf{q} + \rho(\mathbf{z} - \mathbf{u}) \\ \mathbf{b} \end{pmatrix}. \tag{15}$$

Solving this linear system gives the optimal $\mathbf{x}$ and $\boldsymbol{\nu}$ (the Lagrange multipliers for the equality constraint). The ADMM iteration consists of very simple updates to the relevant variables, but its success crucially relies in being able to solve the $\mathbf{x}$-update efficiently. Given the structure of our problem, it is convenient to use a direct solution using Schur's complement, that is:

$$(\mathbf{A}(\mathbf{P} + \rho\mathbf{I})^{-1}\mathbf{A}^T)\boldsymbol{\nu} = \mathbf{A}(\mathbf{P} + \rho\mathbf{I})^{-1}(-\mathbf{q} + \rho(\mathbf{z} - \mathbf{u})) + \mathbf{b} \tag{16a}$$

$$(\mathbf{P} + \rho\mathbf{I})^{-1}\mathbf{x} = -\mathbf{A}^T\boldsymbol{\nu} - \mathbf{q} + \rho(\mathbf{z} - \mathbf{u}). \tag{16b}$$

Eq. (16a) results from left-multiplying the first equation in (15) by $\mathbf{A}(\mathbf{P}+\rho\mathbf{I})^{-1}$ and using the second equation in (15) to eliminate $\mathbf{x}$. Eq. (16b) results from substituting $\boldsymbol{\nu}$ back in the first equation in (15) and solving it for $\mathbf{x}$.

Thus, the ADMM iteration consists of solving a linear system on the primal variables $\mathbf{x}$, applying a thresholding to get $\mathbf{z}$, and an addition to get $\mathbf{u}$. Convergence of the ADMM iteration (12) to the global minimum of problem (7) in value and to a feasible point is guaranteed for any $\rho > 0$.

3.2 Application to our QP

We now write the ADMM updates for our QP (1), where we identify:

$$\begin{aligned}\mathbf{P} &= 2\lambda \operatorname{diag}(\mathbf{L},\ldots,\mathbf{L}) \text{ of } NK \times NK\\ \mathbf{q} &= -\operatorname{vec}(\mathbf{G}) \text{ of } NK \times 1\\ \mathbf{A} &= (\mathbf{I}_N,\ldots,\mathbf{I}_N) \text{ of } N \times NK\\ \mathbf{b} &= \mathbf{1}_N \text{ of } N \times 1\end{aligned}$$

where $\operatorname{vec}(\cdot)$ concatenates the columns of its argument into a single column vector. Given the structure of these matrices, the solution of the KKT system (15) by using Schur's complement (16) simplifies considerably. The basic reasons are that (1) the matrix $\mathbf{P}$ is block-diagonal with K identical copies of the graph Laplacian $\mathbf{L}$, which is itself usually sparse; and (2) the especially simple form of the equality constraint matrix $\mathbf{A}$. Thus, even though the $\mathbf{x}$-update involves solving a large linear system of NK equations, it is equivalent to solving K systems of N equations where the coefficient matrix is the same for each system and besides is constant and sparse, equal to $2\lambda\mathbf{L}+\rho\mathbf{I}$. In turn, these linear systems may be solved efficiently in one of the two following ways: (1) preferably, by caching the Cholesky factorization of this matrix (using a good permutation to reduce fill-in), if it does not add so much fill that it can be stored; or (2) by using an iterative linear solver such as conjugate gradients, initialized with a warm start, preconditioned, and exiting it before convergence, so as to carry out faster, inexact $\mathbf{x}$-updates.

The final algorithm is as follows, with its variables written in matrix form. The input are the affinity matrices $\mathbf{G}_{N\times K}$ and $\mathbf{W}_{N\times N}$, from which we construct the graph Laplacian $\mathbf{L}_{N\times N}$. We then choose $\rho > 0$ and set

$$\mathbf{h} = -\frac{1}{K}\mathbf{G}\mathbf{1}_K + \frac{\rho}{K}\mathbf{1}_N \qquad \mathbf{R}\mathbf{R}^T = 2\lambda\mathbf{L}+\rho\mathbf{I}.$$

The Cholesky factor $\mathbf{R}$ is used to solve linear system (17b). We then iterate, in order, the following updates until convergence:

$$\boldsymbol{\nu} \leftarrow \frac{\rho}{K}(\mathbf{Y}-\mathbf{U})\mathbf{1}_K - \mathbf{h} \tag{17a}$$

$$\mathbf{Z} \leftarrow (2\lambda\mathbf{L}+\rho\mathbf{I})^{-1}(\rho(\mathbf{Y}-\mathbf{U})+\mathbf{G}-\boldsymbol{\nu}\mathbf{1}_K^T) \tag{17b}$$

$$\mathbf{Y} \leftarrow (\mathbf{Z}+\mathbf{U})_+ \tag{17c}$$

$$\mathbf{U} \leftarrow \mathbf{U}+\mathbf{Z}-\mathbf{Y} \tag{17d}$$

where $\mathbf{Z}_{N\times K}$ are the primal variables, $\mathbf{Y}_{N\times K}$ the auxiliary variables, $\mathbf{U}_{N\times K}$ the Lagrange multiplier estimates for $\mathbf{Y}$, and $\boldsymbol{\nu}_{N\times 1}$ the Lagrange multipliers for the equality constraints. The solution for the linear system in the $\mathbf{Z}$-update may be obtained by using two triangular backsolves if using the Cholesky factor of $2\lambda\mathbf{L}+\rho\mathbf{I}$, or using an iterative method such as conjugate gradients if the Cholesky factor is not available.

The iteration (17) is very simple to implement. It requires no line searches and has only one user parameter, the penalty parameter. The algorithm converges for any positive value of the penalty parameter, but this value does affect the convergence rate.

3.3 Remarks

Theorem 3.1. *At each iterate in the algorithm updates* (17), $\mathbf{Z}\mathbf{1}_K = \mathbf{1}_N$, $\mathbf{U} \le \mathbf{0}$, $\mathbf{Y} \ge \mathbf{0}$ *and* $\mathbf{Y} \circ \mathbf{U} = \mathbf{0}$.

Proof. For $\mathbf{Z}$, substituting eq. (17a) into (17b):

$$\mathbf{Z}\mathbf{1}_K = (2\lambda\mathbf{L}+\rho\mathbf{I})^{-1}(\rho(\mathbf{Y}-\mathbf{U})+\mathbf{G}-\boldsymbol{\nu}\mathbf{1}_K^T)\mathbf{1}_K = (2\lambda\mathbf{L}+\rho\mathbf{I})^{-1}(\rho(\mathbf{Y}-\mathbf{U})\mathbf{1}_K+\mathbf{G}\mathbf{1}_K-K\boldsymbol{\nu} = (2\lambda\mathbf{L}+\rho\mathbf{I})^{-1}\rho\mathbf{1}_N = \mathbf{1}_N$$

where the last step results from $(2\lambda\mathbf{L}+\rho\mathbf{I})\mathbf{1}_N = 2\lambda\mathbf{L}\mathbf{1}_N+\rho\mathbf{1}_N = \rho\mathbf{1}_N$. For $\mathbf{U}$, from eqs. (17c)–(17d) we have that $\mathbf{U} \leftarrow \mathbf{U}+\mathbf{Z}-(\mathbf{U}+\mathbf{Z})_+ = (\mathbf{U}+\mathbf{Z})_- \leq \mathbf{0}$. For $\mathbf{Y}$, $\mathbf{Y} \geq \mathbf{0}$ follows from $\mathbf{Y} = (\mathbf{Z}+\mathbf{U})_+$ in (17c). Finally, $\mathbf{Y} \circ \mathbf{U} = \mathbf{0}$ follows from $\mathbf{Y} = (\mathbf{Z}+\mathbf{U})_+$ and $\mathbf{U} = (\mathbf{U}+\mathbf{Z})_-$. □

Theorem 3.2. *Upon convergence of algorithm* (17)*,* $\mathbf{Z}$ *is a solution with Lagrange multipliers* $\boldsymbol{\pi} = -\boldsymbol{\nu}$ *and* $\mathbf{M} = -\rho\mathbf{U}$*.*

Proof. Let us compare the KKT conditions (4) with the algorithm updates (17) upon convergence, i.e., at a fixed point of the update equations. From th. 3.1 we know that $\mathbf{Z}\mathbf{1}_K = \mathbf{1}_N$ and $\mathbf{U} \leq \mathbf{0} \Rightarrow \mathbf{M} = -\rho\mathbf{U} \geq \mathbf{0}$, which are KKT conditions (4b) and (4d). From eq. (17d) we must have $\mathbf{Z} = \mathbf{Y}$, so from eq. (17c) we have $\mathbf{Z} = (\mathbf{Z}+\mathbf{U})_+ \geq \mathbf{0}$, which is KKT condition (4c). From eqs. (17c)–(17d) we have $\mathbf{Z} = (\mathbf{Z}+\mathbf{U})_+$ and $\mathbf{U} = (\mathbf{Z}+\mathbf{U})_-$, therefore $\mathbf{U} \circ \mathbf{Z} = \mathbf{M} \circ \mathbf{Z} = \mathbf{0}$, which is KKT condition (4e). Finally, from eq. (17b) we have:

$$(2\lambda\mathbf{L}+\rho\mathbf{I})\mathbf{Z} = \rho(\mathbf{Y}-\mathbf{U})+\mathbf{G}-\boldsymbol{\nu}\mathbf{1}_K^T \Rightarrow 2\lambda\mathbf{L}\mathbf{Z}-\mathbf{G}+\rho\mathbf{U}+\boldsymbol{\nu}\mathbf{1}_K^T = 2\lambda\mathbf{L}\mathbf{Z}-\mathbf{G}-\boldsymbol{\pi}\mathbf{1}_K^T-\mathbf{M} = \mathbf{0}$$

which matches KKT condition (4a). The change of sign in the multipliers between the algorithm and the KKT conditions is due to the sign choice in the Lagrangian (adding in eq. (11), subtracting in (4)). □

Remark 3.3. In practice, the algorithm is stopped before convergence, and $\mathbf{Z}$, $\boldsymbol{\pi} = -\boldsymbol{\nu}$ and $\mathbf{M} = -\rho\mathbf{U}$ are estimates for a solution and its Lagrange multipliers, respectively. The estimate $\mathbf{Z}$ may not be feasible, in particular the values z_{nk} need not be in $[0,1]$, since this is only guaranteed upon convergence. If needed, a feasible point may be obtained by projecting each row of $\mathbf{Z}$ onto the simplex (see section 4).

Remark 3.4. If $\mathbf{G} = \mathbf{0}$ (or $\lambda \to \infty$) one solution is given by $\mathbf{Z} = \frac{1}{K}\mathbf{1}_N\mathbf{1}_K^T$, for which the Lagrange multipliers are $\mathbf{M} = \mathbf{0}$ and $\boldsymbol{\pi} = \mathbf{0}$, thus the inequality constraints are inactive and the equality constraints are weakly active. Indeed, that $\mathbf{Z}$ value is also a solution of the unconstrained problem.

3.4 Computational complexity

Each step in (17) is $\mathcal{O}(NK)$ except for the linear system solution in (17b). If $\mathbf{L}$ is sparse, using the Cholesky factor makes this step $\mathcal{O}(NK)$ as well, and adds a one-time setup cost of computing the Cholesky factor (which is also linear in N with sufficiently sparse matrices). Thus, each iteration of the algorithm is cheap. In practice, for good values of ρ, the algorithm quickly approaches the solution in the first iterations and then converges slowly, as is known with ADMM algorithms in general. However, since each iteration is so cheap, we can run a large number of them if high accuracy is needed. As a sample runtime, for a problem with $N = 10\,000$ items and $K = 10$ categories (i.e., $\mathbf{Z}$ has 10^5 parameters) and using a 100-nearest-neighbor graph, the Cholesky factorization takes 0.5 s and each iteration takes 0.15 s in a PC.

For large-scale problems, the slow convergence becomes more problematic, and it is possible that the Cholesky factor may create too much fill even with a good preordering. One can use instead an iterative linear solver, such as preconditioned conjugate gradients. Scaling up the training is a topic for future research.

3.5 Initialization

If the LASS problem is itself a subproblem in a larger problem (as in the Laplacian K-modes clustering algorithm; Wang and Carreira-Perpiñán, 2014b), one should warm-start the iteration of eq. (17) from the values of $\mathbf{Y}$ and $\mathbf{U}$ in the previous outer-loop iteration. Otherwise, we can simply initialize $\mathbf{Y} = \mathbf{U} = \mathbf{0}$, which (substituting in eqs. (17a)–(17b)) gives $\mathbf{Z}_0 = (2\lambda\mathbf{L}+\rho\mathbf{I})^{-1}(\mathbf{G}-\frac{1}{K}\mathbf{G}\mathbf{1}_K\mathbf{1}_K^T)+\frac{1}{K}\mathbf{1}_N\mathbf{1}_K^T$ (where $\mathbf{G}-\frac{1}{K}\mathbf{G}\mathbf{1}_K\mathbf{1}_K^T$ is the matrix $\mathbf{G}$ with centered rows, and $\frac{1}{K}\mathbf{1}_N\mathbf{1}_K^T$ is the simplex barycenter). This initialization is closely related to the projection on the simplex of the unconstrained optimum of the LASS problem, as we show next. Consider first the unconstrained minimization

$$\min_{\mathbf{Z}} f(\mathbf{Z}) = \lambda \operatorname{tr}\left(\mathbf{Z}^T\mathbf{L}\mathbf{Z}\right) - \operatorname{tr}\left(\mathbf{G}^T\mathbf{Z}\right).$$

This problem is in fact unbounded unless $\mathbf{G}^T\mathbf{1} = \mathbf{0}$, because taking $\mathbf{Z} = \mathbf{1}\mathbf{v}^T$ for any $\mathbf{v} \in \mathbb{R}^K$, $\mathbf{v} \neq \mathbf{0}$, since $\mathbf{L}\mathbf{1} = \mathbf{0}$, we have $f(\mathbf{1}\mathbf{v}^T) = (\mathbf{G}^T\mathbf{1})^T\mathbf{v}$, which can be made arbitrarily negative. We could still try to define a $\mathbf{Z}$ from the gradient $\nabla f(\mathbf{Z})^T = 2\lambda\mathbf{L}\mathbf{Z} - \mathbf{G}$, but this involves a linear system on $\mathbf{L}$, whose computational cost defeats the purpose of the initialization. Instead, we can consider the unconstrained minimization

$$\min_{\mathbf{Z}} \frac{\rho}{2} \operatorname{tr}\left(\mathbf{Z}^T\mathbf{Z}\right) + \lambda \operatorname{tr}\left(\mathbf{Z}^T\mathbf{L}\mathbf{Z}\right) - \operatorname{tr}\left(\mathbf{G}^T\mathbf{Z}\right)$$

for $\rho > 0$, which is strongly convex and has a unique minimum $\mathbf{Z}^* = (2\lambda\mathbf{L} + \rho\mathbf{I})^{-1}\mathbf{G}$, which we can compute cheaply if we reuse the Cholesky factor for $2\lambda\mathbf{L}+\rho\mathbf{I}$. Now, we can write the initialization $\mathbf{Z}_0$ (for $\mathbf{Y} = \mathbf{U} = \mathbf{0}$) in terms of $\mathbf{Z}^*$ as $\mathbf{Z}_0 = \mathbf{Z}^* - (\frac{1}{K}\mathbf{Z}^*\mathbf{1}_K)\mathbf{1}_K^T + \frac{1}{K}\mathbf{1}_N\mathbf{1}_K^T$, which means that each row vector of $\mathbf{Z}^*$ is translated along the direction $\mathbf{1}_K$. Since this direction is orthogonal to the simplex, both $\mathbf{Z}^*$ and $\mathbf{Z}_0$ have the same projection on it.

Finally, note that if ρ is large, then $\mathbf{Z}_0 \approx \frac{1}{K}\mathbf{1}_N\mathbf{1}_K^T$ and $\mathbf{Z}^* \approx \mathbf{0}$, both of which project onto the simplex barycenter, independently of the problem data.

3.6 Stopping criterion

We stop when $\|\mathbf{Z}^{(\tau+\Delta)} - \mathbf{Z}^{(\tau)}\|_1 < \texttt{tol}$, i.e., when the change in absolute terms in $\mathbf{Z}$ in the last Δ iterations falls below a set tolerance `tol` (e.g. 10^{-5}). Using an absolute criterion here is equivalent to using a relative one, since $\|\mathbf{z}_n\|_1 = 1$, $n = 1, \ldots, N$. Since our iterations are so cheap, evaluating $\left\|\mathbf{Z}^{(\tau+\Delta)} - \mathbf{Z}^{(\tau)}\right\|_1$ takes a runtime comparable to that of the updates in (17) (except for the $\mathbf{Z}$-update, possibly), so testing the stopping criterion every $\Delta = 10$–100 iterations saves around 10% runtime.

Another possible stopping criterion is to test whether the KKT conditions (4) are satisfied up to a given tolerance, using the Lagrange multipliers' estimates provided by the algorithm at each iterate. Each iterate always satisfies (4b) and (4d), so we only need to check (4a), (4c) and (4e) (if the iterate is interior to the inequalities it will also satisfy (4c) and (4e)). Still, it is faster to check for changes in $\mathbf{Z}$.

Since the iterates $\mathbf{Z}$ in the algorithm need not be feasible, they may be slightly infeasible once the stopping criterion is satisfied. If desired, a feasible $\mathbf{Z}$ can be obtained by projecting each assignment vector $\mathbf{z}_n$ onto the simplex (see section 4).

3.7 Optimal penalty parameter ρ

The speed at which ADMM converges depends on the quadratic penalty parameter ρ (Boyd et al., 2011). We illustrate this with the "2 moons" dataset in fig. 1 ($N = 4\,000$ points, $K = 2$ categories, 5-nearest-neighbor graph, $\lambda = 1$), where we set positive similarity values for one point in each cluster, resulting in each cluster being assigned to a different category, as expected. The problem has $8\,000$ parameters and we ran 10^4 iterations, which took 11 s. Little work exists on how to select ρ so as to achieve fastest convergence. Recently, for QPs, Ghadimi et al. (2013) suggest to use $\rho^* = 2\lambda\sqrt{\sigma_{\min}\sigma_{\max}}$ where $\sigma_{\min}$ and $\sigma_{\max}$ are the smallest (nonzero) and largest eigenvalue of the Laplacian. In fig. 1, $\rho^* \approx 0.2$, and we show the relative error $\|\mathbf{Z} - \mathbf{Z}_{\text{opt}}\|_{\text{F}} \,/\, \|\mathbf{Z}_{\text{opt}}\|_{\text{F}}$ vs number of iterations for different ρ/ρ^* (initial $\mathbf{Z} = \mathbf{0}$, with relative error 1). Asymptotically, the convergence is linear; a model $\epsilon_k = r^k\epsilon_0$ gives $r \approx 0.9977$ for ρ^*. While, in the long term, values close to ρ^* work best, in the short term, smaller values are able to achieve an acceptably low relative error ($\approx 10^{-3}$) in just a few iterations, so an adaptive ρ would be best overall.

3.8 Matlab code

The following Matlab code implements the algorithm, assuming a direct solution of the $\mathbf{Z}$-update linear system.

```
function [Z,Y,U,nu] = lass(L,l,G,r,Y,U,maxit,tol)

[N,K] = size(G); LI = 2*l*L+r*speye(N,N); h = (-sum(G,2)+r)/K; Zold = zeros(N,K);
for i=1:maxit
  nu = (r/K)*sum(Y-U,2) - h;
  Z = LI \ bsxfun(@minus,r*(Y-U)+G,nu);
```

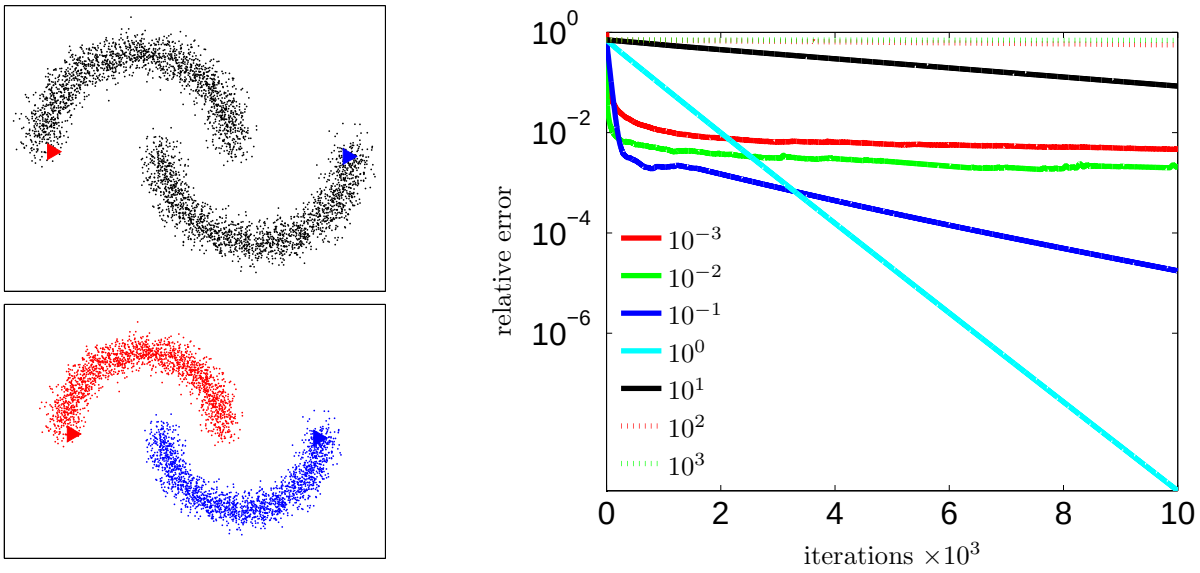

Figure 1: Convergence speed of ADMM for different ρ. *Left*: a 2-moons dataset for which only two points are provided with $\mathbf{g}$ affinities (above) and predicted assignments (below). *Right*: error (in log scale) vs iterations for different ρ/ρ^*.

```
  Y = max(Z+U,0);
  U = U + Z - Y;
  if max(abs(Z(:)-Zold(:))) < tol break; end; Zold = Z;
end
```

4 Out-of-sample mapping

Having trained the system, that is, having found the optimal assignments $\mathbf{Z}$ for the training set items, we are given a new, test item $\mathbf{x}$ (for example, a new point $\mathbf{x} \in \mathbb{R}^D$), along with its item-item and item-category similarities $\mathbf{w} = (w_n)$, $n = 1, \dots, N$ and $\mathbf{g} = (g_k)$, $k = 1, \dots, K$, respectively, and we wish to find its assignment to each category. We follow the reasoning of Carreira-Perpiñán and Lu (2007) to derive an out-of-sample mapping. While one could train again the whole system augmented with $\mathbf{x}$, this would be very time-consuming, and the assignments of all points would change (although very slightly). A more practical and still natural way to define an out-of-sample mapping $\mathbf{z}(\mathbf{x})$ is to solve a problem of the form (1) with a dataset consisting of the original training set augmented with $\mathbf{x}$, but keeping $\mathbf{Z}$ fixed to the values obtained during training. Hence, the only free parameter is the assignment vector $\mathbf{z}$ for the new point $\mathbf{x}$. After dropping constant terms, the optimization problem (1) reduces to the following quadratic program over K variables:

$$\min_{\mathbf{z}} \quad \|\mathbf{z} - (\bar{\mathbf{z}} + \gamma \mathbf{g})\|^2 \tag{18a}$$

$$\text{s.t.} \quad \mathbf{z}^T \mathbf{1}_K = 1 \tag{18b}$$

$$\mathbf{z} \geq \mathbf{0} \tag{18c}$$

where $\gamma = 1/2\lambda(\mathbf{1}_N^T \mathbf{w}) = 1/2\lambda \sum_{n=1}^N w_n$ and

$$\bar{\mathbf{z}} = \frac{\mathbf{Z}^T \mathbf{w}}{\mathbf{1}_N^T \mathbf{w}} = \sum_{n=1}^N \frac{w_n}{\sum_{n'=1}^N w_{n'}} \mathbf{z}_n$$

is a weighted average of the training points' assignments, and so $\bar{\mathbf{z}} + \gamma \mathbf{g}$ is itself an average between this and the item-category affinities. Thus, the solution is the Euclidean projection $\Pi(\bar{\mathbf{z}} + \gamma \mathbf{g})$ of the K-dimensional

vector $\bar{\mathbf{z}}+\gamma\mathbf{g}$ onto the probability simplex. This can be efficiently computed, in a finite number of steps, with a simple $\mathcal{O}(K \log K)$ algorithm (Duchi et al., 2008; Wang and Carreira-Perpiñán, 2014a). Computationally, assuming $\mathbf{w}$ is sparse, the most expensive step is finding the neighbors to construct $\mathbf{w}$. With large N, one should use some form of hashing (Shakhnarovich et al., 2006) to retrieve approximate neighbors quickly.

The out-of-sample prediction for a point in the training set does not generally equal the $\mathbf{z}$ value it received during training (although it does not differ much from it either). That is, $\mathbf{z}(\mathbf{x}_n) \neq \mathbf{z}_n$, where $\mathbf{z}(\mathbf{x}_n)$ uses the training data $\mathbf{w} = \mathbf{W}_{n\cdot}$ and $\mathbf{g} = \mathbf{G}_{n\cdot}$ for $\mathbf{x}_n$. This is also true of semisupervised learning, and it simply reflects the fact that the out-of-sample mapping smoothes, rather than interpolates, the training data.

Given a solution $\mathbf{Z}$ of the LASS training problem, the out-of-sample mapping is uniquely defined, because the problem (18) is strongly convex. However, as described in section 2.2, in particular settings the solution of the LASS training problem may not be unique, and a natural question is: what is the relation between the out-of-sample mappings for two different solutions? From th. 2.1, the solutions have the form $\mathbf{Z}+\mathbf{1}_N\mathbf{p}^T$ where $\mathbf{Z}$ is any particular solution and $\mathbf{p} \in \mathbb{R}^K$ satisfies $\mathbf{p}^T\mathbf{1}_K = 0$, $\mathbf{p}^T(\mathbf{G}^T\mathbf{1}_N) = 0$ and $\mathbf{Z} + \mathbf{1}_N\mathbf{p}^T \geq \mathbf{0}$. Then the out-of-sample mapping for a $\mathbf{p}$-solution has the form

$$\mathbf{z}_\mathbf{p}(\mathbf{x}) = \Pi\left(\frac{(\mathbf{Z}+\mathbf{1}_N\mathbf{p}^T)^T\mathbf{w}}{\mathbf{1}_N^T\mathbf{w}} + \gamma\mathbf{g}\right) = \Pi(\mathbf{v}(\mathbf{x}) + \mathbf{p}) \qquad \mathbf{v}(\mathbf{x}) = \frac{\mathbf{Z}^T\mathbf{w}}{\mathbf{1}_N^T\mathbf{w}} + \gamma\mathbf{g}$$

where $\mathbf{z_0}(\mathbf{x}) = \Pi(\mathbf{v}(\mathbf{x}))$ is the out-of-sample mapping for the base solution $\mathbf{Z}$. If $\mathbf{p}$ was parallel to the vector $\mathbf{1}_K$ then the out-of-sample mappings for different solutions but actually coincide, but in fact $\mathbf{p}^T\mathbf{1}_K = 0$, so the out-of-sample mappings for different solutions correspond to sliding $\mathbf{z_0}(\mathbf{x})$ along the simplex by vector $\mathbf{p}$ (which must respect the remaining conditions above, of course).

As a function of λ, the out-of-sample mapping takes the following extreme values:

- If $\lambda = 0$ or $\mathbf{w} = \mathbf{0}$, $z_k = \delta(k - k_{\max})$ where $k_{\max} = \arg\max\{g_k,\ k = 1, \dots, K\}$, i.e., the item is assigned to its most similar similar category (or any mixture thereof in case of ties).
- If $\lambda = \infty$ or $\mathbf{g} = \mathbf{0}$, $\mathbf{z} = \bar{\mathbf{z}}$, independently of $\mathbf{g}$. This corresponds to the SSL out-of-sample mapping.

In between these, the out-of-sample mapping as a function of λ is a piecewise linear path in the simplex, which represents the tradeoff between the crowd ($\mathbf{w}$) and expert ($\mathbf{g}$) wisdoms. This path is quite different from the simple average of $\bar{\mathbf{z}}$ and $\mathbf{g}$ (which need not even be feasible), and may produce exact 0s or 1s for some entries.

The LASS out-of-sample mapping offers an extra degree of flexibility to the user, which may be used on a case-by-case basis for each test item. The user has the prerogative to set λ to favor more or less the expert vs the crowd opinion, and in fact to explore the entire continuum for $\lambda \in [0, \infty)$. The user can also explore what-if scenarios by changing $\mathbf{g}$ itself, given the vector $\mathbf{w}$ (e.g. how would the assignment vector look like if we think that test item $\mathbf{x}$ belongs to category k but not to category k'?). These computations are all relatively efficient because the bottleneck, which is the $\mathcal{O}(N)$ computation of $\bar{\mathbf{z}}$, is done once only.

Note that the out-of-sample mapping is nonlinear and nonparametric, and it maps an input $\mathbf{x}$ (given its affinity information) onto a valid assignment vector in the probability simplex. Hence, LASS can also be considered as *learning nonparametric conditional distributions over the categories, given partial supervision.*

5 Related work

5.1 Semisupervised learning with a Laplacian penalty (SSL)

In semisupervised learning (SSL) with a Laplacian penalty (Zhu et al., 2003), the basic idea is that we are given an affinity matrix $\mathbf{W}$ and corresponding graph Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{W}$ on N items, and the labels for a subset of the items. Then, the labels for the remaining, unlabeled items are such that they minimize the Laplacian penalty, or equivalently they are the smoothest function on the graph that satisfies the given labels ("harmonic" function). Call $\mathbf{Z}_u$ of $N_u \times K$ and $\mathbf{Z}_l$ of $N_l \times K$ the matrices of labels for the unlabeled and labeled items, respectively, where $N = N_l + N_u$, and $\mathbf{Z}^T = (\mathbf{Z}_u^T\ \mathbf{Z}_l^T)$. To obtain $\mathbf{Z}_u$ we minimize $\operatorname{tr}\left(\mathbf{Z}^T\mathbf{L}\mathbf{Z}\right)$

over $\mathbf{Z}_u$, with fixed $\mathbf{Z}_l$:

$$\min_{\mathbf{Z}_u} \operatorname{tr}\left(\mathbf{Z}^T\mathbf{L}\mathbf{Z}\right) = \min_{\mathbf{Z}_u} \operatorname{tr}\left(\begin{pmatrix}\mathbf{Z}_u\\ \mathbf{Z}_l\end{pmatrix}^T \begin{pmatrix}\mathbf{L}_u & \mathbf{L}_{ul}\\ \mathbf{L}_{ul}^T & \mathbf{L}_l\end{pmatrix}\begin{pmatrix}\mathbf{Z}_u\\ \mathbf{Z}_l\end{pmatrix}\right) = \min_{\mathbf{Z}_u} \operatorname{tr}\left(\mathbf{Z}_u^T\mathbf{L}_u\mathbf{Z}_u + 2\mathbf{Z}_l^T\mathbf{L}_{ul}^T\mathbf{Z}_u\right) + \text{constant}$$

$$\Rightarrow \mathbf{Z}_u = -\mathbf{L}_u^{-1}\mathbf{L}_{ul}\mathbf{Z}_l = \mathbf{L}_u^{-1}\mathbf{W}_{ul}\mathbf{Z}_l. \quad (19)$$

Thus, computationally the solution involves a sparse linear system of $N_u \times N_u$. An out-of-sample mapping for a new test item $\mathbf{x}$ with affinity vector $\mathbf{w}$ wrt the the training set can be derived by SSL again, taking $\mathbf{Z}_l$ of $N \times K$ as all the trained labels (given and predicted) and $\mathbf{Z}_u = \mathbf{z}^T$ as the free label. This gives a closed-form expression

$$\mathbf{z}(\mathbf{x}) = \sum_{n=1}^{N} \frac{w_n}{\sum_{n'=1}^{N} w_{n'}}\mathbf{z}_n \quad (20)$$

which is the average of the labels of $\mathbf{x}$'s neighbors, making clear the smoothing behavior of the Laplacian. SSL with a Laplacian penalty is very effective in problems where there are very few labels, i.e., $N_u \ll N_l$, but the graph structure is highly predictive of each item's labels. Essentially, the given labels are propagated throughout the graph.

In our setting, the labels are the item-category assignments z_{nk}, and we have the following result.

Theorem 5.1. *In problem* (19)*, if* $\mathbf{Z}_l \geq \mathbf{0}$ *and* $\mathbf{Z}_l\mathbf{1}_{N_l} = \mathbf{1}_K$ *then* $\mathbf{Z}_u \geq \mathbf{0}$ *and* $\mathbf{Z}_u\mathbf{1}_{N_u} = \mathbf{1}_K$*.*

Proof. Since $\mathbf{L} = \mathbf{D} - \mathbf{W}$ we have

$$\mathbf{0} = \mathbf{L}\mathbf{1}_{N_u+N_l} = \begin{pmatrix}\mathbf{L}_u & \mathbf{L}_{ul}\\ \mathbf{L}_{ul}^T & \mathbf{L}_l\end{pmatrix}\begin{pmatrix}\mathbf{1}_{N_u}\\ \mathbf{1}_{N_l}\end{pmatrix} = \begin{pmatrix}\mathbf{L}_u\mathbf{1}_{N_u} + \mathbf{L}_{ul}\mathbf{1}_{N_l}\\ \mathbf{L}_{ul}^T\mathbf{1}_{N_u} + \mathbf{L}_l\mathbf{1}_{N_l}\end{pmatrix} \Rightarrow -\mathbf{L}_u^{-1}\mathbf{L}_{ul}\mathbf{1}_{N_l} = \mathbf{1}_{N_u}.$$

Hence $\mathbf{Z}_u\mathbf{1}_K = -\mathbf{L}_u^{-1}\mathbf{L}_{ul}\mathbf{Z}_l\mathbf{1}_K = -\mathbf{L}_u^{-1}\mathbf{L}_{ul}\mathbf{1}_{N_l} = \mathbf{1}_{N_u}$. That $\mathbf{Z}_u \geq \mathbf{0}$ follows from the maximum principle for harmonic functions (Doyle and Snell, 1984): each of the unknowns must lie between the minimum and maximum label values, i.e., in $[0, 1]$. (Strictly, they will lie in $(0, 1)$ or be all equal to a constant.) □

Thus, in the special case where the given labels are valid assignments (nonnegative with unit sum), the predicted labels will also be valid assignments, and we need not subject the problem explicitly to simplex constraints, which simplifies it computationally. This occurs in the standard semisupervised classification setting where each item belongs to only one category and we use the $\mathbf{z}_n$ vectors to implement a 1-of-K coding (e.g. as used for supervised clustering in Grady, 2006). However, *in general SSL does not produce valid assignments*, e.g. if the given labels are not valid assignments, or in other widely used variations of SSL, such as using class mass normalization (Zhu et al., 2003), or using the normalized graph Laplacian instead of the unnormalized one, or using label penalties (Zhou et al., 2004). In the latter case (also similar to the "dongle" variation of SSL; Zhu et al., 2003), one minimizes the Laplacian penalty plus a term equal to the squared distance of the labeled points (considered free parameters as well) to the labels $\mathbf{Z}_l$ provided. Thus, this penalizes the labeled points from deviating from their intended labels, rather than forcing them to equal them. This was extended by Subramanya and Bilmes (2011) (replacing squared losses with Kullback-Leibler divergences and adding an additional entropy term) to learning probability distributions, i.e., where the labels $\mathbf{Z}_l$ are entire distributions over the K classes, with each item-class probability specified exactly. All these approaches rely on the following: *they use provided, specific label values* $\mathbf{Z}_l$ *as targets to be (ideally) met by the parameters.*

Relation with LASS LASS and SSL are similar in that (1) $\mathbf{L}$ plays the same role, i.e., to propagate label information in a smooth way according to the item-item graph; and (2) both rely on some given data to learn $\mathbf{Z}$: the similarity matrix $\mathbf{G}$ in LASS and the given labels $\mathbf{Z}_l$ in SSL. LASS and SSL differ as follows. (1) The use of the simplex constraints, necessary to ensure valid assignments, which also means all the assignment values in LASS are interdependent, unlike in the classical SSL, where the prediction for each category can be solved independently. (2) A fundamental difference is in the supervision provided. If in LASS we were given

actual labels $\mathbf{Z}_l$ for some of the items, we would simply use them just as in SSL, and the LASS problem with $\mathbf{Z}_l$ having given assignments would be:

$$\min_{\mathbf{Z}_u} \quad \lambda \operatorname{tr}\left(\mathbf{Z}_u^T \mathbf{L}_u \mathbf{Z}_u + 2\mathbf{Z}_l^T \mathbf{L}_{ul}^T \mathbf{Z}_u\right) - \operatorname{tr}\left(\mathbf{G}_u^T \mathbf{Z}_u\right) \tag{21a}$$

$$\text{s.t.} \quad \mathbf{Z}_u \mathbf{1}_K = \mathbf{1}_{N_u} \tag{21b}$$

$$\mathbf{Z}_u \geq \mathbf{0}. \tag{21c}$$

It is clear that setting $\mathbf{G} = \mathbf{0}$ there gives the SSL problem (19) if $\mathbf{Z}_l$ contains valid assignments (so the constraints are redundant, from th. 5.1). Thus, SSL is a particular case of LASS, not a different way of encoding the same input data. However, the $\mathbf{G}$ term provides soft, partial "labels", and this information differs from (hard) labels $\mathbf{Z}_l$. Indeed, when the label to be learned for each item is an assignment vector, the concept of "labeling" breaks down, for two reasons. First, if the number of categories K is not very small and an item n has nonzero assignments to multiple categories, in practice it is hard for a user to have to give a value (or tag) for every single relevant category. Giving partial information is much easier, by simply setting $g_{nk} = 1$ for the most relevant categories, possibly setting $g_{nk} = -1$ for a few categories, and setting $g_{nk} = 0$ for the rest (we stick to ± 1 and 0 similarities for simplicity). Second, because the assignment values are constrained to be in the simplex, *we cannot give actual values for individual entries* (unless we give the entire assignment vector). For example, setting an entry to 1 implicitly forces the other entries to 0. In summary, *the semantics of the item-category similarities in LASS is that, where nonzero, they encourage the corresponding assignment towards relatively high or low values (for positive and negative similarities, respectively), and where zero, they reflect ignorance and are non-committing*, something which is close to a user's intuition, but generally difficult to achieve by setting assignment values directly.

Not being able to commit to specific assignment values, especially where $g_{nk} = 0$, also implies that it is not possible to transform meaningfully a given item-category sparse similarity matrix $\mathbf{G}$ into an assignment vector. Given a sparse matrix $\mathbf{G}$, if we insist in setting full assignments for each item n having a nonzero vector $\mathbf{g}_n$ (so we can use these with SSL), perhaps the best one can do is to follow this labeling procedure: for each k, set $z_{nk} = 1$, ϵ or 0 if $g_{nk} = 1$, 0 or -1, respectively, and normalize $\mathbf{z}_n$ (where $1 \gg \epsilon \geq 0$ is a smoothing user parameter). This forces a zero assignment for each negative-similarity category, and distributes the unit assignment mass over the remaining categories. Obviously, this likely forces many z_{nk} to wrong values and, as we show in the experiments, works poorly. (Another approach would be to set only the z_{nk} entries for which $g_{nk} \neq 0$ and leave the others free (subject to the constraints (1b)–(1c)), but this is essentially with LASS does.)

The difference between SSL and LASS is clearly seen in the out-of-sample mapping. In SSL, the information provided for a test item is just the vector $\mathbf{w}$ of similarities to other items, and the SSL out-of-sample mapping coincides with $\bar{\mathbf{z}}$ in the LASS out-of-sample mapping, i.e., the average of its neighbors' assignments. With LASS, in addition to $\mathbf{w}$ we also give the vector $\mathbf{g}$ of similarities to categories. If $\mathbf{g} = \mathbf{0}$, the predictions of LASS and SSL coincide. Otherwise, LASS trades off both $\mathbf{w}$ and $\mathbf{g}$. This is particularly important when $\mathbf{w}$ is not very informative, e.g. if it has many nonzero entries of similar magnitude, or all entries are very small (an outlying or "new" item).

A further difference between SSL and LASS is that in SSL the learned assignments z_{nk} are never exactly 0 or 1 (in nontrivial problems) because each $\mathbf{z}_n$ is the average of its neighbors' assignments. In LASS, the z_{nk} values can be exactly 0 or 1. This happens for the same reason why in many statistical models the L_1 norm achieves sparse solutions while the L_2 norm does not (indeed, the simplex constraints (1b)–(1c) imply $\|\mathbf{z}_n\|_1 = 1$ for each n).

5.2 Assignments and probabilities

The semantics of assignments is different from that of probabilities. Given a discrete probability distribution p over the values in $\Omega = \{1, \ldots, K\}$, saying that "$p(Z = k) = z_k$" means that Z equals exactly one of the values in Ω, but that we observe the value k with a probability z_k (frequentist or Bayesian), for each $k \in \Omega$. In contrast, given an assignment vector $\mathbf{z}$ over the values in Ω, the value of z_k corresponds to the proportion of $\mathbf{z}$ that belongs to category k. For example, in a portfolio model, $k = 1, \ldots, K$ are possible investments and z_k is the portion of a \$1 capital that is allocated to investment k, not the probability that we allocate all the \$1 to investment k. However, in some cases the assignments can indeed be used as proxies for probabilities.

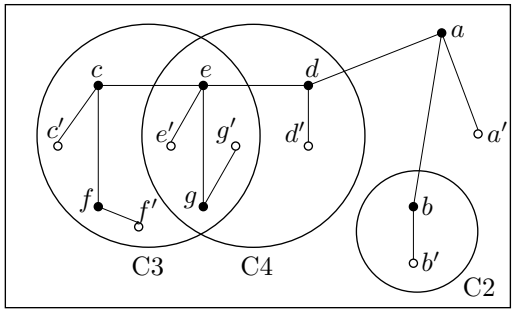

data	ground truth	positive labels	negative labels
a	1	1	not 2
b	1,2	2	
c	1,3	1	
d	1,4	4	not 3
e	1,3,4	1,3	
f	1,3	3	
g	1,3,4	3,4	

SSL

	C1	C2	C3	C4
a'	1	0	0	0
b'	0	1	0	0
c'	1	0	0	0
d'	0	0	0	1
e'	0.5	0	0.5	0
f'	0	0	1	0
g'	0	0	0.5	0.5

LASS (positive labels only)

	C1	C2	C3	C4
a'	0.57	0.28	0.10	0.05
b'	0.40	0.60	0	0
c'	0.30	0	0.70	0
d'	0.36	0.08	0.33	0.23
e'	0.29	0	0.68	0.03
f'	0.05	0	0.95	0
g'	0	0	0.82	0.18

LASS (positive & negative labels)

	C1	C2	C3	C4
a'	0.90	0	0	0.10
b'	0.70	0.30	0	0
c'	0.60	0	0.40	0
d'	0.70	0	0	0.30
e'	0.57	0	0.37	0.07
f'	0.35	0	0.65	0
g'	0.24	0	0.53	0.23

Figure 2: Illustrative example. *Top*: Venn diagram of the categories and ground-truth labels for the labeled items. *Bottom*: assignments found by SSL and LASS (using only positive labels, and both positive and negative labels).

For example, our model includes classification as a particular (and rather restricted) setting, where (most) items belong to a single category and categories are mutually exclusive. Here, the assignment vectors may be interpreted as probabilities and the labels safely provided as 1–of–K codes.

Assignment problems have a long history in operations research and economics. For example, the Markowitz portfolio model (Nocedal and Wright, 2006) seeks the portfolio (soft assignment of an individual investor's \$1 capital to K investments) that maximizes the expected return and minimizes the variance, and has the form of a QP. However, the use of Laplacian penalties (which, in the Markowitz model, would describe the similarity between different investors) does not seem to have been applied there.

5.3 Other applications of LASS

The LASS problem (1) appears in the assignment step of the training algorithm for Laplacian K-modes clustering (Wang and Carreira-Perpiñán, 2014b). Here, the items are data points $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^D$ to be clustered into K clusters and the categories are the cluster centroids $\mathbf{c}_1, \dots, \mathbf{c}_K \in \mathbb{R}^D$. Both affinities are Gaussian: the $\mathbf{W}$ matrix has entries $w_{nm} = G(\|(\mathbf{x}_n - \mathbf{x}_m)/\sigma\|^2)$, $n, m = 1, \dots, N$, and the $\mathbf{G}$ matrix has entries $g_{nk} = G(\|(\mathbf{x}_n - \mathbf{c}_k)/\sigma\|^2)$, $n = 1, \dots, N$, $k = 1, \dots, K$, where $G(\cdot^2)$ gives a Gaussian kernel. At an optimum Laplacian K-modes clustering, the cluster centroids $\mathbf{C} = (\mathbf{c}_1, \dots, \mathbf{c}_K)$ are the modes of the kernel density estimates defined by each of the K clusters' points. The problem is optimized over $\mathbf{Z}$ and $\mathbf{C}$ in alternating steps, where the step over $\mathbf{Z}$ for fixed $\mathbf{C}$ has the form (1).

Supervised clustering can be formulated as problem (1), in particular supervised image segmentation. The user marks several pixels for each segment of the image, from which we can define the similarity matrix $\mathbf{G}$ in various ways. The similarity matrix $\mathbf{W}$ is constructed from the image (pixel location and range features).

6 Experiments

Illustrative example We constructed a simple example to show the difference between LASS and SSL, and the role of positive and negative item-category similarities. The data consists of $N = 14$ points a–g and a'–g' and $K = 4$ categories C1–C4 which are related to each other according to the Venn diagram in

fig. 2. Category C1 contains the other three categories and C3 and C4 intersect. This is quite different from the usual classification setting where each point is assigned to only one category. A graph is built on the data set, denoted by the edges between data points. We consider a–g (filled circles) as partially labeled and a'–g' (hollow circles) as completely unlabeled. The ground-truth (true assignments derived from the Venn diagram), positive and negative labels we use are given in fig. 2(top right). For LASS, we set $g_{nk} = \pm 1$ if $\mathbf{x}_n$ is given positive/negative label Ck, and set $g_{nk} = 0$ if no information is given. We do not give the full label information for any point (except a), so it is crucial to make use of the graph Laplacian to propagate label information. Since each unlabeled point $\mathbf{x}_i$ is only connected to one (partially) labeled point $\mathbf{x}_j$, $\mathbf{x}_i$ will inherit exactly the same assignment from $\mathbf{x}_j$ in order to minimize the objective function (in this case $w_{ij} \left\|\mathbf{z}_i - \mathbf{z}_j\right\|^2 = 0$). Fig. 2(bottom) shows the assignment of unlabeled points obtained from SSL and from LASS ($\lambda = 1$) with only positive labels, and both positive and negative labels in $\mathbf{G}$. With SSL (Zhu et al., 2003), since the labels are used as constraints (eq. (19)), the unlabeled points just inherit them and miss many assignments. With LASS, using only positive $\mathbf{G}$ obtains smoother assignments, but it sometimes assigns points to categories they do not belong to (e.g. a' to C2, d' to C3). Once we use negative labels, the progagation of those wrong labels is cut off and we obtain assignments closer to the ground truth for all points. SSL cannot easily use the negative labels.

Digit recognition We test LASS in a classification task where each data point has only one valid label. We randomly sample 10 000 MNIST digit images and compute the 10-nearest-neighbor graph with similarities based on the Euclidean distance between images. We then randomly select N_l images from each of the 0–9 categories and give them the correct label. We compare with a nearest neighbor classifier (NN), one-vs-all kernel support vector machine (KSVM) using RBF kernel of width $\sigma = 5$ and hinge loss penalty parameter C selected from $\{10^{-3}, 10^{-2}, 10^{-1}, 1, 10^1, 10^2, 10^3\}$ and SSL (Zhu et al., 2003; Grady, 2006) using 1-out-of-10 coding for the labeled points. We also include two variants of SSL: SSL1 (Zhu et al., 2003) normalizes the assignments from SSL so that the prior distribution of the different classes is respected. SSL2 uses the normalized graph Laplacian instead of the unnormalized one in SSL (Zhou et al., 2004). SSL1 and SSL2 improve over SSL but neither of them produce valid assignments (they do not lie on the probability simplex). For LASS and SSL/1/2 we assign each point to the category with the highest prediction value. We let all algorithms use their respective optimal parameters (e.g. λ in LASS is determined by a grid search).

Fig. 3(left) shows the classification error over 20 different labeled/unlabeled partitions of the dataset as a function of N_l (errorbars not shown to avoid clutter). The accuracy of all algorithms improves as the number of labeled points increase, particularly for NN and KSVM, which are template matchers. But when only few points are labeled, the methods that make use of Laplacian smoothing significantly outperform them. LASS (runtime: 40 s) consistently achieves the best accuracy while producing valid assignments.

Document categorization We want to predict assignments of documents to topics, where each document may belong to multiple topics, in the 20-newsgroups dataset ($N = 11\,269$ documents). We manually add 7 new topics (comp.sys, rec.sport, computer, recreation, politics, science and religion) based on the hierarchical structure and perceived similarity of groups (e.g. comp.sys.ibm.pc.hardware / comp.sys.mac.hardware). This yields $K = 27$ topics and each document can belong to 1 to 3 topics. To construct feature vectors, we remove words that occur in 5 or fewer documents, and then extract the TFIDF (term frequency $\times$ inverse document frequency) feature of the documents. We generate the $\mathbf{G}$ similarity matrix by randomly selecting N_l documents from each of the 27 topics, and giving each document one $+1$ label (the topic it is selected from) and five -1 labels (topics it does not belong to). For SSL, we turn this into assignment labels as described in section 19, and select the smoothing parameter ϵ optimally from {0, 0.005, 0.01, 0.02, 0.05, 0.1, 0.2}.

To evaluate the performance, we select for each test document the T topics to which it has highest predicted assignments (where $T \in \{1, 2, 3\}$ is the actual number of topics this document belongs to), consider them as predicted label set, and compare them with ground truth labels. We consider it as an error if the predicted label set and the ground truth set differ. NN classification does not apply here because no document is fully labeled. We can apply one-vs-all linear SVM (hinge loss penalty parameter C selected from $\{10^{-3}, 10^{-2}, 10^{-1}, 1, 10^1, 10^2, 10^3\}$) because we do have training points for each topic. Fig. 3(right) shows the mean classification error over 20 random labeled/unlabeled partitions of the dataset as a function of N_l. The accuracy of all algorithms again improves when N_l increases. LASS outperforms all other algorithms at nearly all N_l values, and, unlike them, always produces valid assignments.

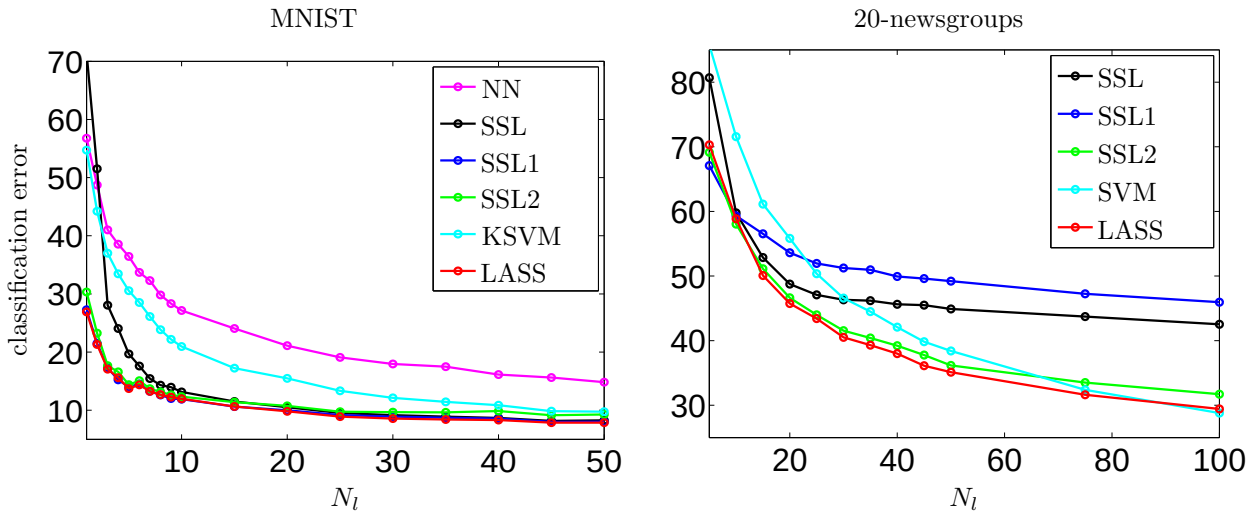

Figure 3: Classification error (%) vs number of labeled points for each class on MNIST (left) and 20-newsgroups (right) datasets.

Image tagging This experiment fully benefits from the ability of LASS to handle partial labels, and predict full assignment vectors. In the problem of image tagging, each image can typically be tagged with multiple categories of a large number of possible categories. However, a user will usually tag only a few of the relevant categories for it (e.g. out of laziness) and will miss tagging other categories that would be relevant too. The task given a test image is to predict an assignment vector, i.e., to fill in "soft tags", for all categories. When there are many possible categories, trying to fill in the missing assignments for the partially labeled samples (so we can use SSL to propagate them to the unlabeled samples) is pointless. In contrast, LASS does not require these missing assignments, by conveniently providing zero affinities.

We demonstrate LASS on a subset of the ESP game (von Ahn and Dabbish, 2004) images used by Guillaumin et al. (2009). We select the images in the training set that are tagged with at least 6 categories (words), resulting in 6 100 images with a total of 267 non-empty categories, with 7.2 categories per image on average. We use the same image feature sets as Guillaumin et al. (2009) to compute distances between images and build a 10-nearest neighbor graph. We give partial information for 4 600 images and provide item-category affinities for each image in the following way: we give positive affinity (+1) for a random subset of size n_l from the categories it is tagged with, and give negative affinities (−1) randomly for 5 out of the 20 most frequent categories it is not tagged with. Providing negative affinities in this way stops the algorithm from concentrating most of the probability mass on the most frequent categories. The other 1 500 images are completely unlabeled and used for testing. SSL fills in missing assignments of the partially labeled samples as described in the document categorization experiment. Parameters are selected based on grid search for each algorithm.

We evaluate the performance of different algorithms using the precision, recall and F–1 score (averaged over sample images) on the test samples while fixing the annotation length at 5, i.e., each image is tagged with the 5 categories of highest assignment. (Although LASS admits tags as similarity values at test time, we do not use them here.) We vary the number of positive tags n_l from 1 to 6. Fig. 4 shows the results for SSL and LASS over 20 runs (each with a different random selection of test set and partial affinities). We could not run one-versus-all SVMs because there are no negative samples for most categories. In SSL2, the highest prediction values are nearly always the most frequent categories. We see that LASS greatly improves over SSL, especially when smaller numbers of positive affinities are given.

7 Conclusion

We have proposed a simple quadratic programming model for learning assignments of items to categories that combines two complementary and possibly conflicting sources of information: the crowd wisdom and the

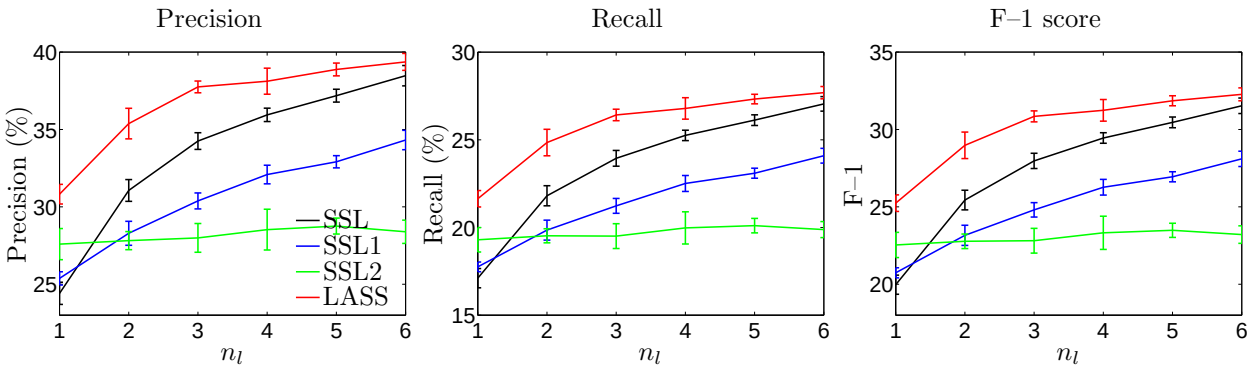

GT: dog grass green man sky white
Pred.: grass man sky green white **tree**

GT: black drawing man old soldier tent white
Pred.: black old drawing white tent **sketch** man

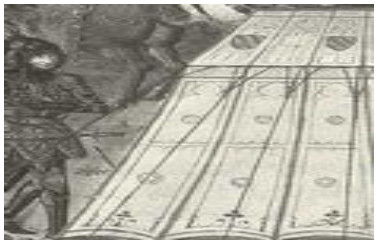

GT: blue computer gray purple screen window
Pred.: computer screen gray window blue **white**

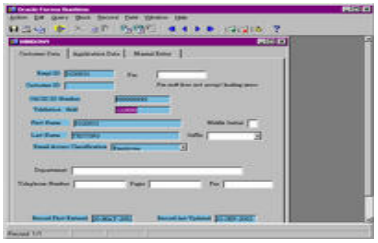

GT: black drawing hair man nose old white
Pred.: black white drawing man hair **circle tie**

GT: black coin man money old round silver white
Pred.: black old round coin money man **woman gray**

GT: field grass green people sky tree
Pred.: grass sky green **man** tree **tent**

Figure 4: Results on ESP game dataset. *Top*: precision and recall vs number of positive affinities n_l, at a fixed annotation length 5, averaged over 20 runs. *Bottom*: sample predictions of LASS on test images when $n_l = 5$. We show the highest predictions (*Pred.*, in assignment order) up to the number of tags in the ground truth (*GT*, unordered), with **mismatches** in boldface.

expert wisdom. It is particularly attractive when fully labeling an item is impractical, or when categories have a complex structure and items can genuinely belong to multiple categories to different extents. It provides a different way to incorporate supervision to that of traditional Laplacian semisupervised learning, which is ill-suited for this setting because the similarity information cannot be faithfully transformed into assignment labels.

We have derived a training algorithm based on the alternating direction method of multipliers. The algorithm's iterations can be made fast by caching the Cholesky factorization of the graph Laplacian. The algorithm is very simple to implement. It requires no line searches and has only one user parameter, the penalty parameter. The algorithm converges for any positive value of the penalty parameter, but this value does affect the convergence rate.

We expect LASS to apply to problems beyond semisupervised learning, such as clustering, and in social

network applications, with image, sound or text data that is partially tagged. It can also be extended to handle tensor data or have additional terms in its objective, for example to represent relations between categories with a category-category similarity matrix, or even to use negative similarities in the item-item graph (since the feasible set is bounded, solution(s) still exist). A further application of LASS is to learn probability distributions that are conditional on partial supervisory information, since, in effect, the out-of-sample mapping is a nonparametric mapping from the affinity information to a distribution over categories. Another research direction is to find optimally adaptive schedules for the penalty parameter and to accelerate the convergence of the training algorithm, particularly with large datasets with many categories, where we may expect each row of $\mathbf{Z}$ to be sparse.

References

M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, June 2003.

M. Belkin, P. Niyogi, and V. Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *J. Machine Learning Research*, 7:2399–2434, Nov. 2006.

S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1): 1–122, 2011.

M. Á. Carreira-Perpiñán. Fast nonparametric clustering with Gaussian blurring mean-shift. In W. W. Cohen and A. Moore, editors, *Proc. of the 23rd Int. Conf. Machine Learning (ICML'06)*, pages 153–160, Pittsburgh, PA, June 25–29 2006.

M. Á. Carreira-Perpiñán and Z. Lu. The Laplacian Eigenmaps Latent Variable Model. In M. Meilă and X. Shen, editors, *Proc. of the 11th Int. Workshop on Artificial Intelligence and Statistics (AISTATS 2007)*, pages 59–66, San Juan, Puerto Rico, Mar. 21–24 2007.

M. Á. Carreira-Perpiñán and W. Wang. LASS: A simple assignment model with Laplacian smoothing. In C. E. Brodley and P. Stone, editors, *Proc. of the 28th National Conference on Artificial Intelligence (AAAI 2014)*, Quebec City, Canada, July 27–31 2014.

O. Chapelle, B. Schölkopf, and A. Zien, editors. *Semi-Supervised Learning*. Adaptive Computation and Machine Learning Series. MIT Press, 2006.

F. R. K. Chung. *Spectral Graph Theory*. Number 92 in CBMS Regional Conference Series in Mathematics. American Mathematical Society, Providence, RI, 1997.

P. G. Doyle and J. L. Snell. *Random Walks and Electric Networks*, volume 22 of *Carus Mathematical Monographs*. Mathematical Association of America, 1984.

J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the ℓ_1-ball for learning in high dimensions. In A. McCallum and S. Roweis, editors, *Proc. of the 25th Int. Conf. Machine Learning (ICML'08)*, pages 272–279, Helsinki, Finland, July 5–9 2008.

E. Ghadimi, A. Teixeira, I. Shames, and M. Johansson. Optimal parameter selection for the alternating direction method of multipliers (ADMM): Quadratic problems. arXiv:1306.2454 [math.OC], July 3 2013.

L. Grady. Random walks for image segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(11):1768–1783, Nov. 2006.

M. Guillaumin, T. Mensink, J. Verbeek, and C. Schmid. TagProp: Discriminative metric learning in nearest neighbor models for image auto-annotation. In *Proc. 12th Int. Conf. Computer Vision (ICCV'09)*, pages 309–316, Kyoto, Japan, Sept. 29 – Oct. 2 2009.

Z. Lu and M. Á. Carreira-Perpiñán. Constrained spectral clustering through affinity propagation. In *Proc. of the 2008 IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR'08)*, Anchorage, AK, June 23–28 2008.

J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer-Verlag, New York, second edition, 2006.

G. Shakhnarovich, P. Indyk, and T. Darrell, editors. *Nearest-Neighbor Methods in Learning and Vision*. Neural Information Processing Series. MIT Press, Cambridge, MA, 2006.

J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(8):888–905, Aug. 2000.

A. Subramanya and J. Bilmes. Semi-supervised learning with measure propagation. *J. Machine Learning Research*, 12:3311–3370, 2011.

G. Taubin. A signal processing approach to fair surface design. In S. G. Mair and R. Cook, editors, *Proc. of the 22nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH 1995)*, pages 351–358, Los Angeles, CA, Aug. 6–11 1995.

L. von Ahn and L. Dabbish. Labeling images with a computer game. In *Proc. ACM Int. Conf. Human Factors in Computing Systems (CHI 2004)*, pages 319–326, Vienna, Austria, Apr. 24–29 2004.

W. Wang and M. Á. Carreira-Perpiñán. Manifold blurring mean shift algorithms for manifold denoising. In *Proc. of the 2010 IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR'10)*, pages 1759–1766, San Francisco, CA, June 13–18 2010.

W. Wang and M. Á. Carreira-Perpiñán. The role of dimensionality reduction in classification. In C. E. Brodley and P. Stone, editors, *Proc. of the 28th National Conference on Artificial Intelligence (AAAI 2014)*, Quebec City, Canada, July 27–31 2014a.

W. Wang and M. Á. Carreira-Perpiñán. Laplacian K-modes clustering. Unpublished manuscript, 2014b.

D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf. Learning with local and global consistency. In S. Thrun, L. K. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems (NIPS)*, volume 16, pages 321–328. MIT Press, Cambridge, MA, 2004.

X. Zhu, Z. Ghahramani, and J. Lafferty. Semi-supervised learning using Gaussian fields and harmonic functions. In T. Fawcett and N. Mishra, editors, *Proc. of the 20th Int. Conf. Machine Learning (ICML'03)*, pages 912–919, Washington, DC, Aug. 21–24 2003.